Node-Specific Triad Pattern Mining for Complex-Network Analysis

Marco Winkler
Institute for Theoretical Physics
University of Wuerzburg
Am Hubland, 97074 Wuerzburg, Germany
Email: mwinkler@physik.uni-wuerzburg.de

Joerg Reichardt
Institute for Theoretical Physics
University of Wuerzburg
Am Hubland, 97074 Wuerzburg, Germany
Email: reichardt@physik.uni-wuerzburg.de

arXiv:1410.1594v2 [cs.SI] 15 Jan 2015

***Abstract*—The mining of graphs in terms of their local substructure is a well-established methodology to analyze networks. It was hypothesized that motifs - subgraph patterns which appear significantly more often than expected at random - play a key role for the ability of a system to perform its task. Yet the framework commonly used for motif-detection averages over the local environments of all nodes. Therefore, it remains unclear whether motifs are overrepresented in the whole system or only in certain regions.**

In this contribution, we overcome this limitation by mining *node-specific* triad patterns. For every vertex, the abundance of each triad pattern is considered only in triads it participates in. We investigate systems of various fields and find that motifs are distributed highly heterogeneously. In particular we focus on the feed-forward loop motif which has been alleged to play a key role in biological networks.

I. Introduction

The concept of networks has been successfully applied to model the interactions between entities in complex systems of various fields. The topological structure of interactions among the constituents of complex many particle systems is intimately linked to system function and global system properties. A major branch of networks research aims to elucidate this link between structure and function.

Generally, pairwise relations between nodes, so-called dyads, are considered the fundamental building blocks of complex networks. Erdös-Renyi graphs [1], the configuration model [2], stochastic block models [3], [4], [5] and degree corrected block models [6] all fall into the class of dyadic models. The basic assumption underlying dyadic models is that dyads are *conditionally independent* given the model's parameters.

However, the assumption of dyadic independence as a general paradigm of network modeling seems questionable. For example, in a social context, the idea that the relation of Alice and Bob be independent from the relation of Alice and Charlie seems to go against experience, especially if the relation is of romantic type. Similarly, triadic closure, or the large clustering coefficient observed in many networks, hints at a dependence between the connections in a network. Generalizing these ideas, during the last decade the systematic study of higher order sub-network structure has attracted much attention [7], [8], [9], [10], [11], [12]. Subgraph patterns which occur significantly more often than expected at random, so called *motifs*, have been hypothesized to serve as the acutal building blocks of many network topologies [7].

Published in IEEE ICDMW 2014. ©2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Fig. 1. All 13 possible connected non-isomorphic triadic subgraphs (subgraph patterns) in directed unweighted networks.

The framework commonly applied to analyze local network structure detects an *average* over- or underrepresentation of subgraph patterns for the whole system. Yet, especially for larger complex networks, there may be areas in which one structural pattern is of importance, whereas in different regions other patterns are relevant. The contribution of this work is to overcome this limitation by introducing a framework for *node-specific* pattern mining. More specifically, instead of mining frequent subgraph patterns of the whole system, we investigate the neighborhood of every single node separately. I.e. for every node we consider only triads it participates in. In this paper, we further apply our new tool to multiple real-world data sets by analyzing their node-specific triad patterns. We show that for many systems their motifs are distributed highly heterogenously.

The remainder of this work will be organized as follows. In Section II we briefly review work related to the analysis of networks in terms of their local substructure. In Section III we describe our methodology of node-specific pattern mining. Results obtained from experiments on real datasets will be presented in Section IV.

II. Related Work

In the context of local subgraph analysis, most attention has been devoted to the investigation of triadic subgraphs [7], [13], [8], [9], [10]. Apart from node permutations, there are 13 distinct connected triad patterns in directed unweighted

networks as shown in Figure 1. It was found that certain patterns of third-order subgraphs occur significantly more frequently than expected in an ensemble of networks with the same degree distributions as the original network.

Over- and underrepresentation of pattern i in a graph is quantified through a Z-score

$$Z_i = \frac{N_{\text{original},i} - \langle N_{\text{rand},i} \rangle}{\sigma_{\text{rand},i}}. \quad (1)$$

$N_{\text{original},i}$ is the number of appearances of pattern i over all possible 3-tuples of nodes in the original network. Sampling from the ensemble of randomized networks yields the average occurrence $\langle N_{\text{rand},i} \rangle$ of that pattern and the respective standard deviation $\sigma_{\text{rand},i}$. Thus, the Z-scores represent a measure for the significance of an over- or underrepresentation for each pattern i shown in Figure 1.

Hence, every network can be assigned a vector $\vec{Z}$ whose components comprise the Z-scores of all possible triad patterns. Significant patterns are referred to as 'motifs' [7]. Further, one commonly refers to the normalized Z-vector as the 'significance profile', $\vec{SP} = \vec{Z}/\sqrt{\sum_i Z_i^2}$. This normalization makes systems of different sizes comparable [7].

A multitude of real-world systems has been examined in terms of their triadic Z-score profiles [8], [14], [15], [16] and it was found that systems with similar tasks tend to have similar Z-score profiles and thus exhibit the same motifs. Therefore, it was conjectured that their structural evolution may have been determined by the relevance of these motif patterns. Hence, motifs, rather than independent links, have been suspected to serve as the basic entities of complex network structure [7]. In particular, the role of the 'feed-forward loop' pattern ( ) has been discussed intensively in the field [13], [17], [18], [19]. It has been alleged to play a key role for systems to reliably perform information-processing tasks.

However, there has also been ongoing discussion about the expressive power of the subgraph-analysis described above [20]. E.g. latent-variable models might offer an explanation for many of the observed motif distributions. The randomization employed in the subgraph-mining process ignores all mesoscopic structure, potentially present in the system. Hence, parts of the non-vanishing Z-scores may stem from such structure [21], [22]. Comparison of a network with block structure to a null model which does not account for such groups may result in over- or underrepresentations of patterns which are more than less artifacts of the mesoscopic structure. It was further shown that there are pairwise correlations in the Z-scores of the subgraph patterns in Figure 1 which occur solely for statistical reasons and therefore have to be taken into account when interpreting the functional role of motif patterns [23], [24].

In the present work we aim to further unravel the role of triad motifs in complex networks. We introduce a methodology for node-specific pattern mining which allows us to localize the regions of a graph in which the instances of a motif predominantly appear. Thus, it is possible to identify and remove the nodes and links which eventually make a certain pattern a motif. This will enable future investigations to assess whether it is actually the presence of a motif which enables a system to perform its task or whether other structural aspects are more relevant.

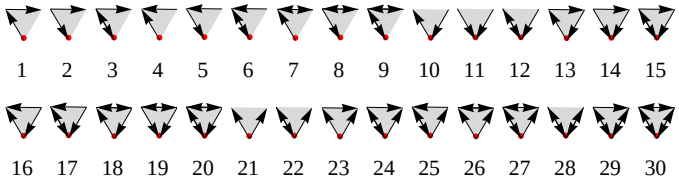

Fig. 2. All possible connected, nonisomorphic triadic subgraph patterns in terms of a distinct node (here: lower node).

III. Node-Specific Triad Pattern Mining (NoSPaM)

Based on the commonly applied subgraph-detection procedure we will now introduce the algorithm for **No**de-**S**pecific **Pa**ttern **M**ining (NoSPaM). For every node α in a graph NoSPaM evaluates the abundance of all structural patterns in α's neighborhood, i.e. patterns α participates in, and compares it to the expected frequency of occurrence in a randomized ensemble of the network under investigation. In the latter, both individual in- and out degrees of all nodes, and the number of unidirectional and bidirectional links are the same as in the original network.

The NoSPaM framework can be realized for patterns comprised of an arbitrary number of nodes n. The detailed algorithmic implementation of NoSPaM$_n$ is then dependent on the pattern size. For the rest of this work we will focus on triad patterns ($n = 3$). However, a generalization to higher orders is straightforward. Although, of course, with increasing n, the number of non-isomorphic subgraphs increases rapidly.

We aim to evaluate the abundance of triad patterns from a node α's point of view. Therefore, the symmetry of most patterns of Figure 1 is now broken and the number of connected *node-specific* triad patterns increases to 30 rather than 13. They are shown in Figure 2. To understand the increase in the number of patterns consider the oridinary subgraph . From the perspective of one particular node, this pattern splits into the three node-specific triad patterns 1, 5, and 10 in Figure 2.

Of course, some patterns are included in others, e.g. pattern 1 in pattern 3. It shall be emphasized that, in order to avoid biased results, we do not double count, i.e. an observation of pattern 3 will only increase its corresponding count and not the one associated with pattern 1.

For every node α in a graph, NoSPaM$_3$ will now compute Z-scores for each of the 30 node-specific patterns i shown above:

$$Z_i^\alpha = \frac{N^\alpha_{\text{original},i} - \left\langle N^\alpha_{\text{rand},i} \right\rangle}{\sigma^\alpha_{\text{rand},i}}. \quad (2)$$

$N^\alpha_{\text{original},i}$ now is the number of appearances of pattern i in the triads where node α participates in. Accordingly, $\left\langle N^\alpha_{\text{rand},i} \right\rangle$ is the expected frequency of pattern i in the triads where

node α is part of in the randomized ensemble. $\sigma^{\alpha}_{\text{rand},i}$ is the corresponding standard deviation.

The algorithm for the node-specific triad pattern mining will now be presented in three parts. Algorithm 1 describes the randomization process which is the same as the shuffling done for ordinary subgraph mining. The microscopic steps performed to randomize a network are illustrated in Figure 3.

Algorithm 1 Degree-preserving randomization of a graph

```
function RANDOMIZE(Graph G, no. of required steps)
    success = 0
    while success < number of required rewiring steps do
        pick a random link e1 ∈ G
        if e1 is unidirectional then
            pick a 2nd unidirectional link e2 ∈ G at random
        else
            pick a 2nd bidirectional link e2 ∈ G at random
        end if
        if e1 and e2 do not share a node then
            rewire according to the rules in Figure 3(a)
            success++
            if one of the new links already exists then
                undo the rewiring
                success−−
            end if
        end if
    end while
    return randomized instance of G
end function
```

The number of rewiring steps should be chosen proportionally to the number of links in the graph. Since all randomization steps preserve individual node degrees and the number of unidirectional and bidirectional links, these quantities are also conserved on the macroscopic level. It shall be mentioned that alternative randomization methods exist: exponential random graph models (ERGMs) allow for a faster generation of randomized networks [25], however they come with the limitation to fix only the expectation for individual node degrees, not necessarily the actual values. For applications of the methodology to big data ERGMs may serve as an alternative to generate the random null models, yet going along with a loss of accuracy.

Algorithm 2 performs the counting process for the appearances of triad patterns in a graph. Because it is computationally expensive to test all triads in the system ($\mathcal{O}\left(N^3\right)$), we rather iterate over pairs of adjacent edges in the graph. Since real-world networks are usually sparse, this is much more efficient.

Using the functions defined in Algorithm 1 and Algorithm 2, we can eventually formulate the routine of our new methodology of node-specific triad pattern mining (NoSPAM$_3$). Algorithm 3 describes its formalism. It computes the node-specific Z-scores as defined in Equation 2. All operations on arrays in Algorithm 3 are performed elementwise.

Performance: The computational cost of Algorithm 1, C_1, scales with the number of required randomization steps per

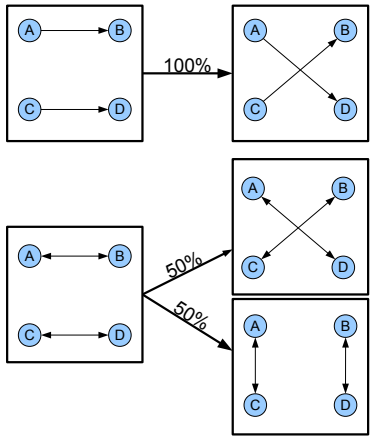

(a)

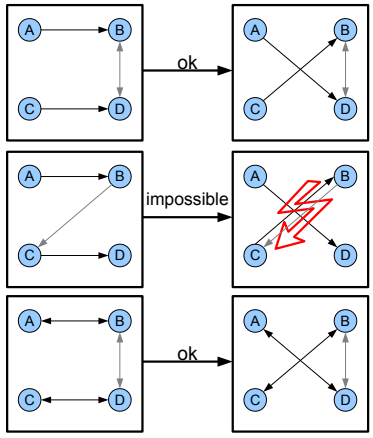

(b)

Fig. 3. Microscopic link-switchings performed to generate the randomized ensembles. (a) illustrates the switching rules. (b) shows cases where additional links affect the randomization process: only if these links interfer with the rewiring of a particular edge (middle) the step cannot be performed.

Algorithm 2 Counting of node-specific triad patterns

```
function NSPPATTERNCOUNTER(Graph G)
    N: N × 30-dimensional array storing the pattern counts
        for every node of G
    for every edge e ∈ G do
        i, j ← IDs of e's nodes with i < j
        C ← {} be list of candidate nodes to form triad
                patterns including e
        C ← all neighbors of i
        C ← all neighbors of j
        for all c ∈ C do
            if i + j < sum of IDs of all other connected
                        dyads in triad (ijc) then
                increase the counts in N for i, j, and c for
                their respective node-specific patterns
            end if
        end for
    end for
    return N
end function
```

Algorithm 3 Node-specific triad pattern mining (NOSPAM$_3$)

function NOSPAM(Graph $\mathcal{G}$, # required rewiring steps, # randomized instances)
  $\mathcal{N}_{\text{original}} \leftarrow$ NSPPATTERNCOUNTER($\mathcal{G}$)
  $\mathcal{N}_{\text{rand}} \leftarrow \{\}$
  $\mathcal{N}_{\text{sq,rand}} \leftarrow \{\}$
  for # randomized instances **do**
    $\mathcal{G} \leftarrow$ RANDOMIZE($\mathcal{G}$, # required rewiring steps)
    counts $\leftarrow$ NSPPATTERNCOUNTER($\mathcal{G}$)
    $\mathcal{N}_{\text{rand}} \leftarrow \mathcal{N}_{\text{rand}} +$ counts
    $\mathcal{N}_{\text{sq,rand}} \leftarrow \mathcal{N}_{\text{sq,rand}} +$ counts $*$ counts
  end for
  $\mathcal{N}_{\text{rand}} \leftarrow \mathcal{N}_{\text{rand}}/(\#\text{randomized instances})$
  $\mathcal{N}_{\text{sq,rand}} \leftarrow \mathcal{N}_{\text{sq,rand}}/(\#\text{randomized instances})$
  $\sigma_{\text{rand}} \leftarrow \sqrt{\mathcal{N}_{\text{sq,rand}} - (\mathcal{N}_{\text{rand}} * \mathcal{N}_{\text{rand}})}$
  $\mathcal{Z} \leftarrow (\mathcal{N}_{\text{original}} - \mathcal{N}_{\text{rand}})/\sigma_{\text{rand}}$
  return $\mathcal{Z}$
end function

instance, which should be chosen proportionally to the number of edges E in the graph $\mathcal{G}$, i.e. $C_1 = \mathcal{O}(E)$.

Algorithm 2 iterates over all edges of $\mathcal{G}$ and their adjacent edges. Therefore, it is $C_2 = \mathcal{O}(E \cdot k_{\max}) \leq \mathcal{O}(E^2)$ where $k_{\max}$ is the maximum node degree in $\mathcal{G}$. In real-world networks $k_{\max}$ is usually much smaller than E.

Finally, the total computational cost of NOSPAM$_3$ (Algorithm 3) depends on the desired number of randomized network instances I. Algorithm 2 is invoked $(1 + I)$ times, Algorithm 1 is invoked I times. Hence, the total computational cost is

$$C_{\text{NOSPAM}_3} = \mathcal{O}(I \cdot E \cdot k_{\max}). \quad (3)$$

Further, it shall be mentioned that NOSPAM$_3$ is parallelizeable straightforwardly since the evaluations in terms of the randomized network instances can be executed independently of each other.

IV. RESULTS

We will now present results obtained from the application of NOSPAM$_3$ to peer-reviewed real-world datasets. All networks are directed and edges are treated as unweighted. An implementation of the pattern-mining program is made publically available online [26].

Yeast transcriptional [27], [28]**:** 688 nodes, 1,079 edges. Transcriptional network of the yeast S. Cerevisiae. Nodes are genes, edges point from regulating genes to regulated genes. It is not distinguished between activation and repression.

E. Coli transcriptional [27], [17]**:** 423 nodes, 519 edges. Nodes are operons, each edge is directed from an operon that encodes a transcription factor to an operon that it directly regulates (an operon is one or more genes transcribed on the same mRNA).

Neural network of C. Elegans [29], [30]**:** 279 nodes, 2,194 edges. Nodes are the neurons of the largest connected component in the somatic nervous system. Edges describe the chemical synapses between the neurons.

Scientific citations [31], [32]**:** 27,700 nodes, 352,807 edges. Nodes are high-energy physics papers on the arXiv, submitted between January 1993 and April 2003. Edges from node A to B indicate that paper A cites paper B. Although it may seem unintuitive, there are papers citing each other.

Political blogs [33], [34]**:** 1,224 nodes, 19,025 edges. Largest connected component of a network where the nodes are political blogs. Edges represent links between the blogs recorded over a period of two months preceding the 2004 US Presidential election.

French book [27], [8]**:** 8,325 nodes, 24,295 edges. Word-adjacency network of a French book. Nodes are words, an edge from node A to B indicates that word B directly follows word A at least once in the text.

Spanish book [27], [8]**:** 11,586 nodes, 45,129 edges. Word-adjacency network of a Spanish book.

A. Node-specific vs. ordinary triadic Z-score profiles

Figure 4 shows the node-specific triadic Z-score profiles for various systems. We used at least five switches per edge (on average) for every iteration and averaged over 1000 iterations. Note that there is one curve for every node in the graph. The node-specific patterns on the horizontal axis are oriented the way that the node under consideration is the lower one.

We find that systems from similar fields have similar node-specific triadic Z-score profiles. Figures 4(a) and 4(b) show transcriptional networks, Figures 4(c) and 4(d) show data from a social context, specifically the citation network of high-energy physicists and the network of hyperlinks between political blogs. Figures 4(e) and 4(f) show word-adjacency networks in French and Spanish language, respectively. The observation that systems from a similar context exhibit similar local structural aspects indicate that the latter are strongly linked to the systems' function.

The fact that NOSPAM$_3$ provides localized data enables us to identify the areas of a graph where certain subgraph patterns primarily occur. Particularly, it allows us to test whether motifs of a system are overabundant throughout the entire network or if they are restricted to limited regions or the proximity of few nodes. In order to explore this issue, for each node, we will map its node-specific Z-scores to a score for the regular triad patterns (shown in Figure 1). This will be realized by taking the mean over the Z-scores of all node-specific triad patterns corresponding to a regular triad pattern. The mapping is shown in Table I. Hence, we obtain a 13-dimensional mapped node-specific Z-score profile for every node in a graph.

The gray, thin curves in Figure 5 show the mapped scores for each node for multiple real-world networks. In addition, the red, thick curve shows the regular Z-score profile over the whole network obtained by the well-established motif-detection analysis of Milo et al. [8]. Although the gray and the red curves are not independent of each other, it shall be mentioned that the regular Z-score profile can not be computed

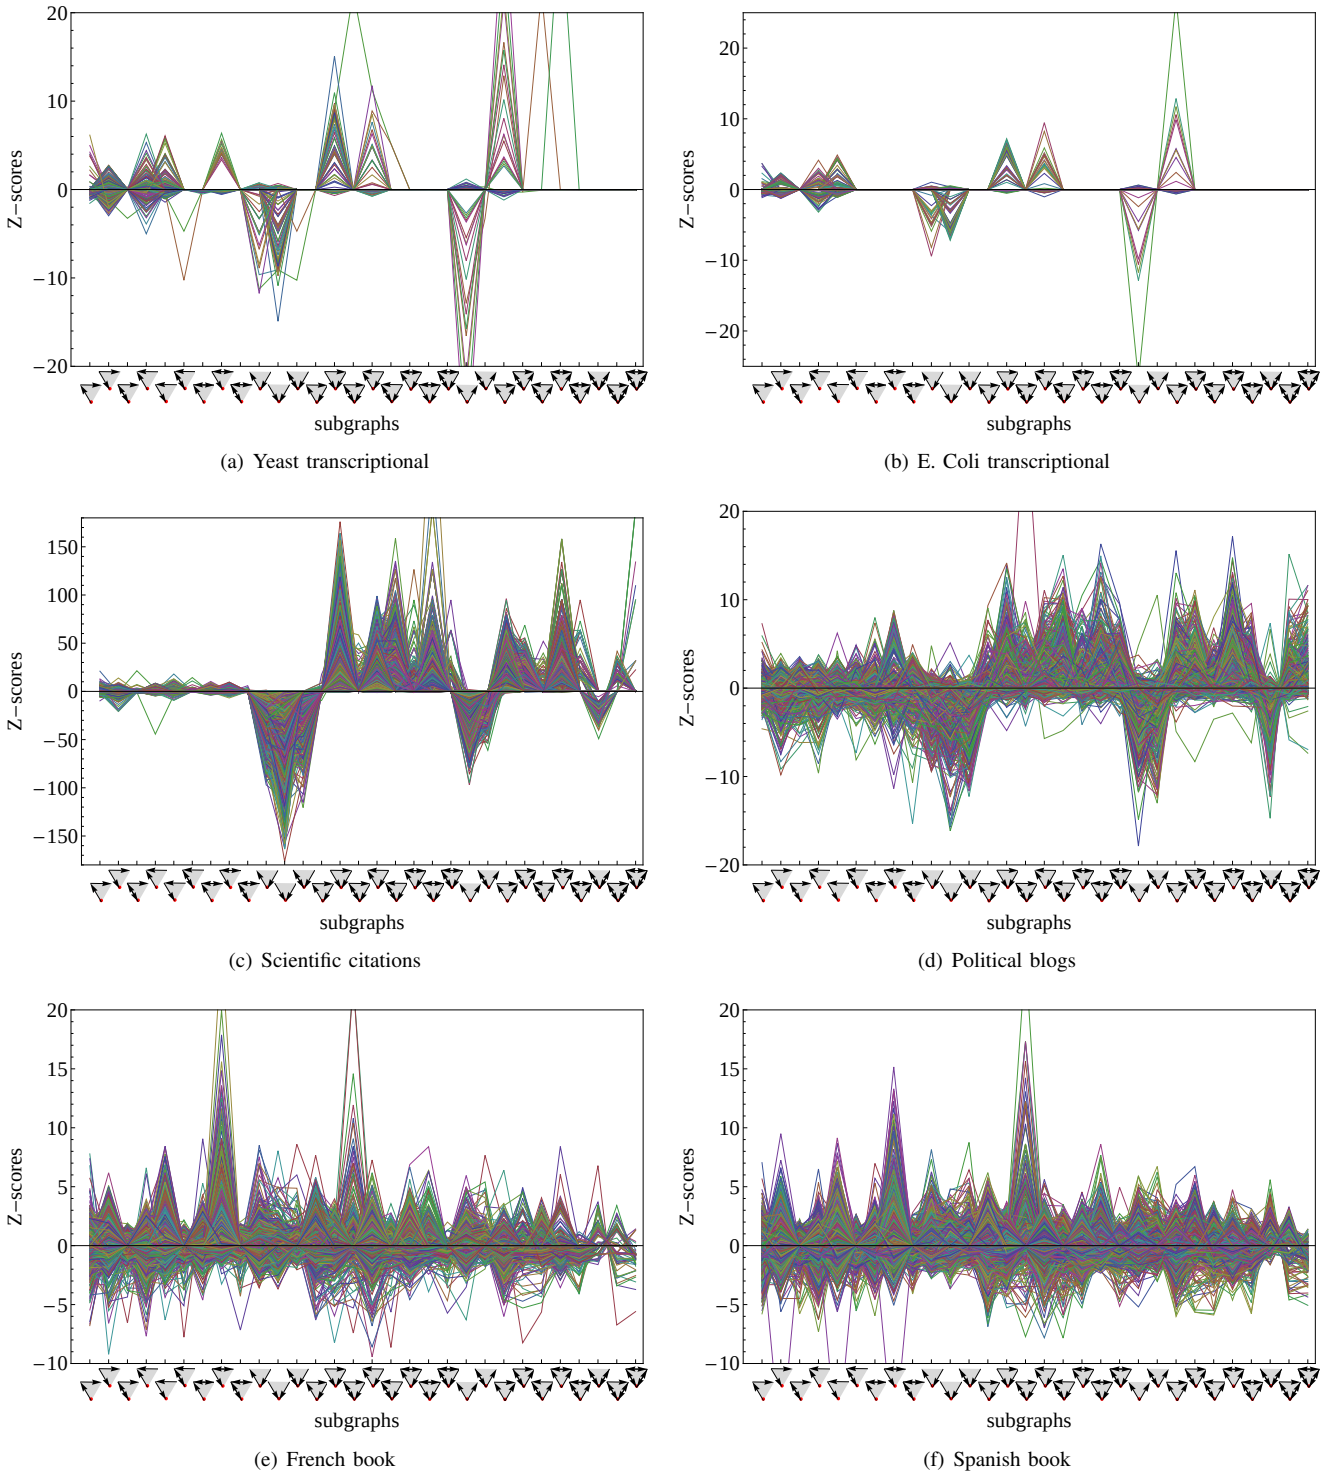

(a) Yeast transcriptional

(b) E. Coli transcriptional

(c) Scientific citations

(d) Political blogs

(e) French book

(f) Spanish book

Fig. 4. Node-specific Z-score profiles of various real-world networks: transcriptional networks of a) the yeast S. Cerevisiae [27], [28] and b) E. Coli [27], [17], socially related networks such as c) the citations of scientific papers [31], [32] and d) hyperlinks between political blogs [33], [34], and word-adjacency networks of e) French books [27], [8] and f) Spanish books [27], [8]. The node-specific patterns on the horizontal axis are oriented the way that the node under consideration is the lower one.

TABLE I

MAPPING OF NODE-SPECIFIC TRIAD PATTERNS TO THEIR REGULAR TRIAD PATTERNS.

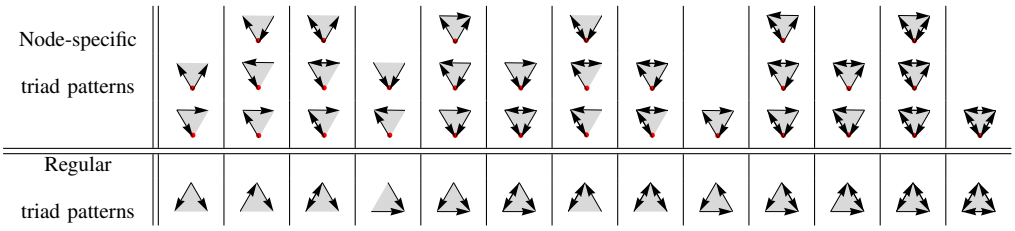

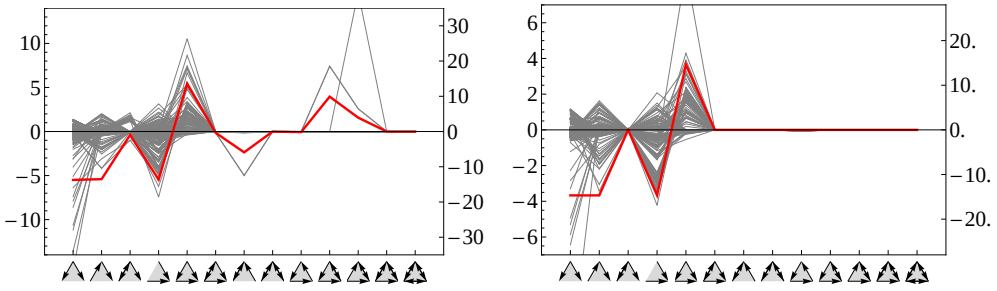

(a) Yeast transcriptional

(b) E. Coli transcriptional

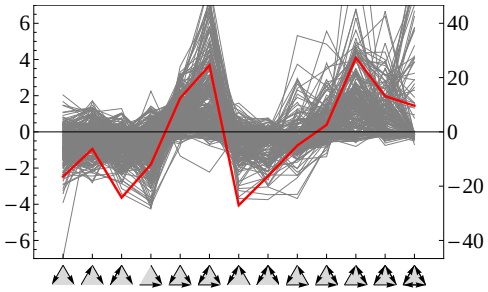

(c) C.Elegans neural network

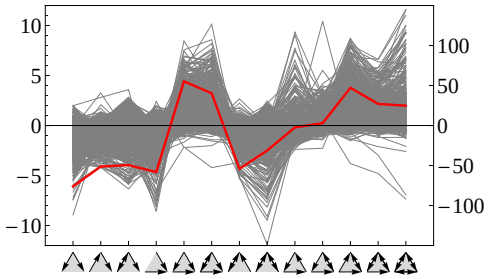

(d) Political blogs

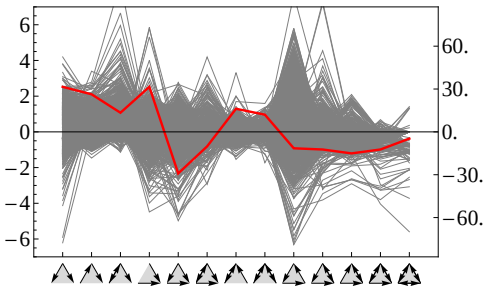

(e) French book

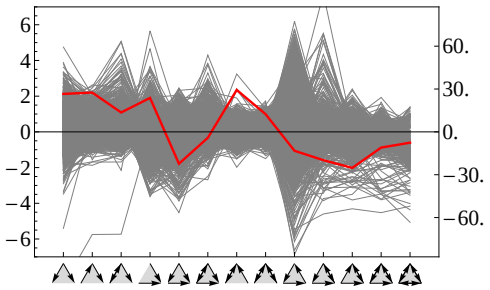

(f) Spanish book

Fig. 5. Node-specific triadic Z-scores mapped to the patterns of Figure 1. For each pattern the average is taken over all corresponding node-specific patterns (Table I). The scaling on the left corresponds to the node-specific triad patterns, the one on the right to the Z-scores of the ordinary triad patterns.

from the gray curves directly. In particular it is not the mean of the latters.

It can be observed that even though a pattern may be overrepresented referring to the system as a whole, it may still be underrepresented in the neighborhood of certain nodes. Moreover, there are patterns with a rather low regular Z-score, while there are both nodes with a strong positive and nodes with a strong negative contribution to the pattern. These contradictory effects seem to compensate each other on the system level. The decribed phenomenon can be particularly observed in the word-adjacency networks in Figures 5(e) and 5(f), especially for the loop pattern ( ).

B. Heterogeneity of motif distributions

To further investigate whether motifs appear homogenously distributed over a graph we will devote ourselves to the feed-forward loop (FFL) pattern ( ). The FFL is one of the patterns which has been studied most intensively with respect to its relevance for systems to reliably perform their functions [13], [17], [18], [19]. Specifically in transcriptional regulation networks it was argued that the FFL pattern might play an important role for facilitating its information-processing tasks [13].

Figure 6 shows two of those transcriptional regulation networks. In both of them the FFL is a motif. Vertex sizes are scaled by the magnitude of the averaged node-specific Z-scores of the three patterns corresponding to the FFL. Positive contributions are shown with bright vertices, negative ones are filled (do only occur with very small magnitude, i.e. very small node sizes). Apparently there are no nodes in the networks with a significant negative contribution to the FFL. Yet, neither is the pattern homogenously overrepresented throughout the whole system, although it is a motif. In fact, for most nodes the FFL-subgraph structure does not seems to play any role whatsoever. In contrast, there are few nodes with a rather strong contribution to the FFL eventually making it a motif of the entire system.

This effect becomes even clearer when considering histograms over the nodes' FFL contributions of the two systems. Figure 7(a) shows the histogram of S. Cerevisiae, Figure 7(b) the one of E. Coli. Both exhibit a strong peak around zero indicating that most nodes do not participate in FFL structures significantly. Only very few nodes have a large mean node-specific triadic Z-score for the patterns corresponding to the FFL.

There are two potential implications which can be derived from these observations: One could be that the FFL motif is actually not that important for the systems to work reliably. The second consequence could be that in fact very few nodes are critical for the system to work the way it is supposed to do. In the second case system would be very prone to failure of these crucial vertices. It may be subject to future research to further investigate these possible implications for dynamical processes on different topologies and under node failure.

V. Conclusions

With this work we have introduced a novel tool for the analysis of complex networks in terms of their local substructure. Existing methods have focussed on the detection of patterns, which appear more frequently than expected for a random null model, over the network *as a whole*. However, it has remained unclear whether these network motifs are overrepresented homogenously over the systems or whether they are bound to the neighborhood of certain nodes in the networks.

To overcome this limitation we introduced the framework of node-specific pattern mining, NoSPaM. Rather than averaging over the local topology of the entire system, for every node in the graph, we evaluate Z-score profiles which describe the nodes' individual local topologies.

For the analyis of real-world data sets we applied NoSPaM_3 which analyzes the local topology in terms of triadic subgraph patterns. We found that systems of similar fields tend to have similar node-specific triadic Z-score profiles indicating that local structural aspects are intimately connected with the systems' function.

Considering the mean contribution of node-specific triadic Z-scores to their respective ordinary triad patterns we found that the appearance of certain subgraphs is distributed highly heterougenously for many systems. This observation was supported by investigating the appearance of the feed-forward loop (FFL) pattern in more detail for transcriptional regulation networks in which it is a motif. The functional relevance of the FFL for systems to properly perform their task has been discussed intensively in complex networks research. It was conjectured that the FFL plays a key role for systems to process information. The fact that it appears highly heterogenously distributed over the graphs raises the question about its actual role for the systems. Will their function be significantly shortened when these nodes fail or are motifs in fact not important to such an extent? And how is the evolution of dynamical processes on systems affected by nodes with certain subgraph characteristics? To answer these questions will be subject to future research.

It may be further promising to adapt the analysis to signed networks in order to investigate issues related to social balance research. Moreover, the NoSPaM analysis of a network yields a new set of structural features for each node. Using these properties as inputs for clustering, classification, or role-detection algorithms can give rise to a better understanding of network designs and to improve link prediction on graphs.

References

[1] P. Erdös and A. Rényi, "On random graphs, I," *Publicationes Mathematicae (Debrecen)*, vol. 6, pp. 290–297, 1959.

[2] M. Molloy and B. Reed, "A critical point for random graphs with a given degree sequence," *Random Structures & Algorithms*, vol. 6, no. 2-3, pp. 161–180, Mar. 1995.

[3] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," *Social Networks*, vol. 5, no. 2, pp. 109–137, Jun. 1983.

[4] Y. J. Wang and G. Y. Wong, "Stochastic Blockmodels for Directed Graphs," *Journal of the American Statistical Association*, vol. 82, no. 397, pp. 8–19, 1987.

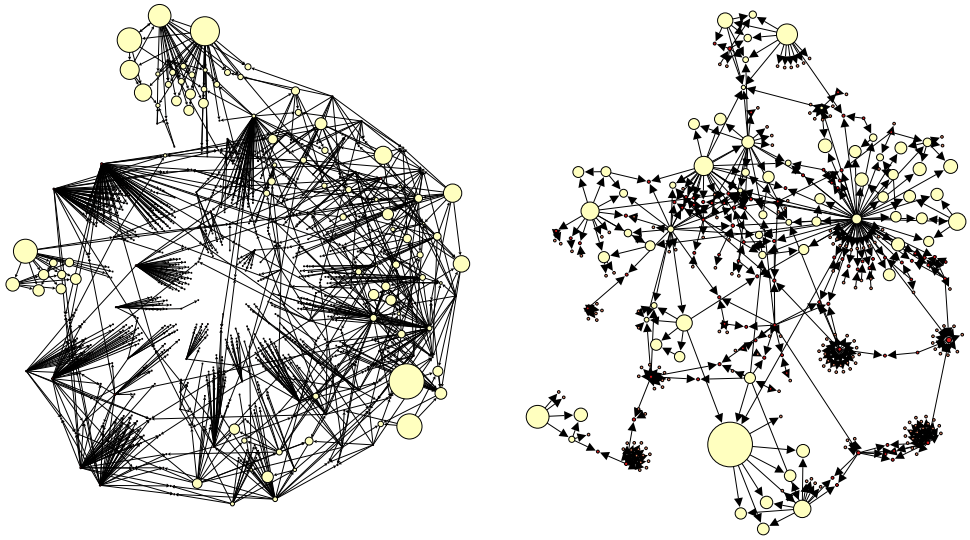

(a) Yeast (b) E. Coli

Fig. 6. Transcriptional regulatory networks. Vertex sizes indicate the magnitude of the mean node-specific triad Z-scores corresponding to the feed-forward loop motif (numbers 14, 16, and 23 in Figure 2). White vertices indicate positive values, filled vertices correspond to negative values (occur only with very small, hardly visible nodes). Graph plots were produced with Gephi [35].

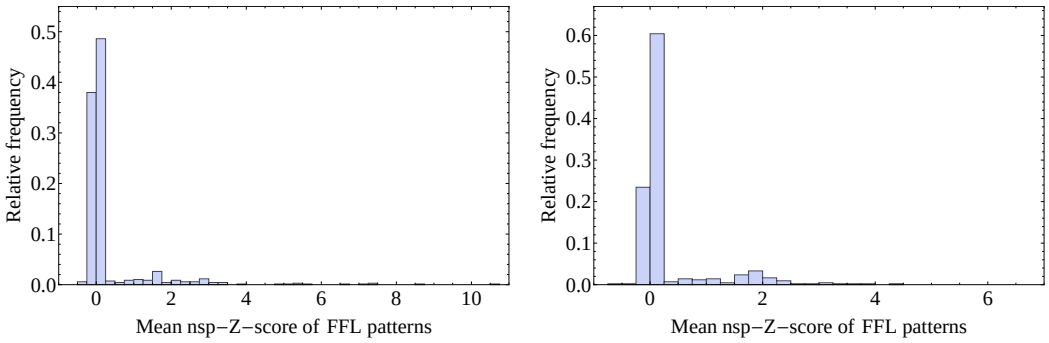

(a) Yeast transcriptional (b) E. Coli transcriptional

Fig. 7. Histograms for the mean node-specific triad Z-scores corresponding to the feed-forward loop motif (numbers 14, 16, and 23 in Figure2).

[5] K. Nowicki and T. A. B. Snijders, "Estimation and Prediction for Stochastic Blockstructures," *Journal of the American Statistical Association*, vol. 96, no. 455, pp. 1077–1087, 2001.

[6] B. Karrer and M. E. J. Newman, "Stochastic blockmodels and community structure in networks." *Physical Review E - Statistical, Nonlinear and Soft Matter Physics*, vol. 83, no. 1 Pt 2, p. 016107, 2010.

[7] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, "Network motifs: simple building blocks of complex networks." *Science*, vol. 298, no. 5594, pp. 824–7, Oct. 2002.

[8] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon, "Superfamilies of evolved and designed networks." *Science*, vol. 303, no. 5663, pp. 1538–42, Mar. 2004.

[9] O. Sporns and R. Kötter, "Motifs in Brain Networks," *PLoS Biology*, vol. 2, no. 11, p. e369, 2004.

[10] I. Albert and R. Albert, "Conserved network motifs allow protein-protein interaction prediction." *Bioinformatics (Oxford, England)*, vol. 20, no. 18, pp. 3346–52, Dec. 2004.

[11] M. Berlingerio, F. Bonchi, B. Bringmann, and A. Gionis, "Mining graph evolution rules," in *Machine learning and knowledge discovery in databases*. Springer, 2009, pp. 115–130.

[12] M. Rahman, M. Bhuiyan, and M. Hasan, "Graft: An efficient graphlet counting method for large graph analysis," 2014.

[13] S. S. Shen-Orr, R. Milo, S. Mangan, and U. Alon, "Network motifs in the transcriptional regulation network of Escherichia coli." *Nature genetics*, vol. 31, no. 1, pp. 64–8, May 2002.

[14] Y.-H. Eom, S. Lee, and H. Jeong, "Exploring local structural organization of metabolic networks using subgraph patterns." *Journal of theoretical biology*, vol. 241, no. 4, pp. 823–9, Aug. 2006.

[15] P. Kaluza, A. Kölzsch, M. T. Gastner, and B. Blasius, "The complex network of global cargo ship movements." *Journal of the Royal Society*, vol. 7, no. 48, pp. 1093–103, Jul. 2010.

[16] S. Sakata, Y. Komatsu, and T. Yamamori, "Local design principles of mammalian cortical networks." *Neuroscience research*, vol. 51, no. 3, pp. 309–15, Mar. 2005.

[17] S. Mangan and U. Alon, "Structure and function of the feed-forward loop network motif." *PNAS*, vol. 100, no. 21, pp. 11 980–5, Oct. 2003.

[18] S. Mangan, A. Zaslaver, and U. Alon, "The coherent feedforward loop serves as a sign-sensitive delay element in transcription networks." *Journal of molecular biology*, vol. 334, no. 2, pp. 197–204, Nov. 2003.

[19] U. Alon, "Network motifs: theory and experimental approaches." *Nature reviews. Genetics*, vol. 8, no. 6, pp. 450–61, Jun. 2007.

[20] Y. Artzy-Randrup, S. J. Fleishman, N. Ben-Tal, and L. Stone, "Comment on "Network motifs: simple building blocks of complex networks" and "Superfamilies of evolved and designed networks"." *Science*, vol. 305, no. 5687, p. 1107; author reply 1107, Aug. 2004.

[21] J. Reichardt, R. Alamino, and D. Saad, "The Interplay between Microscopic and Mesoscopic Structures in Complex Networks," *PLoS ONE*, vol. 6, no. 8, p. e21282, Aug. 2011.

[22] M. E. Beber, C. Fretter, S. Jain, N. Sonnenschein, M. Müller-Hannemann, and M.-T. Hütt, "Artefacts in statistical analyses of network motifs: general framework and application to metabolic networks." *Journal of the Royal Society*, vol. 9, no. 77, pp. 3426–35, Dec. 2012.

[23] R. Ginoza and A. Mugler, "Network motifs come in sets: Correlations in the randomization process," *Physical Review E*, vol. 82, no. 1, pp. 1–6, Jul. 2010.

[24] M. Winkler and J. Reichardt, "Motifs in triadic random graphs based on Steiner triple systems," *Physical Review E*, vol. 88, no. 2, p. 022805, Aug. 2013.

[25] T. De Bie, "Maximum entropy models and subjective interestingness: an application to tiles in binary databases," *Data Mining and Knowledge Discovery*, vol. 23, no. 3, pp. 407–446, 2011.

[26] M. Winkler, "Personal Website," http://www.mwinkler.eu, 2014.

[27] U. Alon, "Uri Alon Lab," http://www.weizmann.ac.il/mcb/UriAlon/, 2012.

[28] M. C. Costanzo, "YPDTM, PombePDTM and WormPDTM: model organism volumes of the BioKnowledgeTM Library, an integrated resource for protein information," *Nucleic Acids Research*, vol. 29, no. 1, pp. 75–79, Jan. 2001.

[29] L. R. Varshney, B. L. Chen, E. Paniagua, D. H. Hall, and D. B. Chklovskii, "Structural properties of the Caenorhabditis elegans neuronal network." *PLoS computational biology*, vol. 7, no. 2, p. e1001066, Jan. 2011.

[30] Z. Alton, L. Herndon, C. Crocker, R. Lints, and D. Hall, "http://www.wormatlas.org," 2002-2012.

[31] J. Gehrke, P. Ginsparg, and J. Kleinberg, "Overview of the 2003 kdd cup," *SIGKDD Explor. Newsl.*, vol. 5, no. 2, pp. 149–151, Dec. 2003.

[32] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graphs over time: Densification laws, shrinking diameters and possible explanations," in *Proc. of the 11th ACM SIGKDD Internat. Conf. on Knowledge Discovery in Data Mining*. New York, NY, USA: ACM, 2005, pp. 177–187.

[33] M. E. J. Newman, "http://www-personal.umich.edu/∼mejn/netdata/," 2012.

[34] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 US election: divided they blog," in *Proceedings of the 3rd international workshop on Link discovery*. ACM, 2005, pp. 36–43.

[35] M. Bastian, S. Heymann, and M. Jacomy, "Gephi: An open source software for exploring and manipulating networks," *International AAAI Conference on Weblogs and Social Media*, 2009.